

PREDICTIVE MOVER DETECTION AND TRACKING IN CLUTTERED ENVIRONMENTS

Luis Navarro-Serment, Christoph Mertz, and Martial Hebert
The Robotics Institute, Carnegie Mellon University
Pittsburgh, PA, 15213

ABSTRACT

This paper describes the design and experimental evaluation of a system that enables a vehicle to detect and track moving objects in real-time. The approach investigated in this work detects objects in LADAR scan lines and tracks these objects (people or vehicles) over time. The system can fuse data from multiple scanners for 360° coverage. The resulting tracks are then used to predict the most likely future trajectories of the detected objects. The predictions are intended to be used by a planner for dynamic object avoidance. The perceptual capabilities of our system form the basis for safe and robust navigation in robotic vehicles, necessary to safeguard soldiers and civilians operating in the vicinity of the robot.

1. INTRODUCTION

Safe navigation is one of the most important goals for any vehicle. To operate in real-world environments, vehicles must successfully avoid collisions with other moving objects (people or vehicles) while traversing the environment.

The ability to avoid colliding with other moving objects is particularly important in autonomous vehicles. This is especially important in cases where the vehicle operates in close proximity with people. In order to be effective, a vehicle's collision avoidance system must perform two basic tasks: detect and track moving objects. The timely detection of an object makes the vehicle aware of a potential danger in its vicinity. Similarly, the vehicle can predict the most likely future positions of an object being tracked, and make corrections to its present course accordingly. For instance, a vehicle tracking a pedestrian currently walking on the sidewalk in the same direction may decide to continue its present course. However, if the vehicle anticipates that a pedestrian walking ahead of it is about to cross the street, it must then either slow down or stop completely.

Robust and reliable detection and tracking has attracted a lot of attention in recent years, driven by applications such as pedestrian protection (Fuerstenberg and Scholz, 2005), vehicle platooning, and autonomous driving (Sun et al., 2006). This is a difficult problem, which becomes even harder when the sensors (e.g., optical sensors, radar, laser scanners) are mounted on the

vehicle rather than being fixed, such as in traffic monitoring systems. Effective detection and tracking require accurate measurements of object position and motion, even when the sensor itself is moving. Range sensors are well suited to this problem because a first-order motion correction can be made by simply subtracting out self-motion from range measurements. Unfortunately, merely subtracting out ego-motion does not eliminate all the effects of motion because the perceived object's shape seems to change as different aspects of the object come into view, and this change can easily be misinterpreted as motion. Plus, the perceived appearance of an object depends on its pose, and can also be affected by nearby objects. Finally, complex outdoor environments frequently involve cluttered backgrounds, unpredictable interaction between traffic participants, and are difficult to control.

The fundamental problem is that, in order to detect change in the object's position, it is necessary to observe some fixed reference point on it. However, if the reference point is not truly fixed, then false apparent motion is perceived. In other words, apparent shape change due to changing perspective can be misinterpreted as motion. The severity of the shape change problem depends primarily on the largest object size, the slowest speed to be measured and the time available for detection. How much can the reference point shift? If the apparent center of the object is used as the reference, then due to angular resolution limits, the reference can shift by more than 1/2 the object size in a short time. This happens when the long side of an object suddenly becomes visible.

In this paper, we describe the design of a system that enables a vehicle to detect and track moving objects in real-time (Fig. 1). The approach investigated detects objects in LADAR scan lines and tracks these objects (people or vehicles) over time. The tracker detects moving objects and estimates their position and motion, while largely ignoring self-motion-induced changes in the scan. The resulting tracks are then used to predict the most likely near-future trajectories of the detected objects and generate collision warnings. Our work differs from previous approaches in that the detection-tracking-prediction elements are integrated into a single system.

The evaluation of tracking systems is difficult, since it is hard to provide target ground truth. A formal assessment of such systems in vehicular applications is

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 01 NOV 2006		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Predictive Mover Detection And Tracking In Cluttered Environments				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The Robotics Institute, Carnegie Mellon University Pittsburgh, PA, 15213				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADM002075., The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			



Figure 1. A Demo III Experimental Unmanned Vehicle (named XUV), and a smaller robot used as a controlled target for establishing ground truth.

rarely found in the literature. Consequently, we also present the experimental evaluation of system performance using a small robot as a controlled-motion target to establish ground truth.

Finally, we present tracking results from controlled experiments using pedestrians, as well as an evaluation of object motion predictions in a collision warning system.

2. RELATED WORK

The problem of detection and tracking of moving objects for vehicular applications has received considerable attention in recent years. The most commonly used approaches involve both active and passive sensors (Hebert, 2000). Active sensors, such as radar and LADAR, detect the distance of objects by measuring the travel time of a signal emitted by the sensor and reflected by the object. Conversely, passive sensors, such as video cameras, acquire data in a non-intrusive way. (Sun et al., 2006) present an extensive review of vision-based on-road vehicle detection systems.

Active sensors have the advantage of being capable of measuring certain quantities (e.g., distance) directly without requiring powerful computing resources. In particular, recent models of laser scanners are capable of gathering high resolution data at high scanning speeds, and are available in enclosures suitable for vehicular applications. The closest work related to our approach involves the use of laser line scanners, and it is described in the rest of this section.

In (Fuerstenberg et al., 2002), the authors describe the application of a multilayered laser scanner for pedestrian classification. Vehicle odometry is used to estimate self-motion, removing the kinematic effects of sensor motion. A Kalman filter is used for object velocity estimation. Tracked objects are classified as car, pedestrian, etc., based on their apparent shape and behavior over time. Fuerstenberg’s work also produced a

second system (Streller et al., 2002), in which a Kalman filter estimates motion based on the change in position of an object’s estimated center-point. Object classification is used to fit a class-specific prior rectangular model to the points. Although not mentioned explicitly, this appears to be an approach to reducing shape-change motion artifacts. The success of this technique would depend on the correctness of the classification and the prior model. Each object class also has distinct fixed Kalman filter parameters. A multi-hypothesis approach is used to mitigate the effect of classification error. The emphasis of both efforts is on single-LADAR systems, and multi-scanner fusion is not considered.

In (Wang et al., 2003) the authors generalize Simultaneous Localization and Mapping (SLAM) to allow detection of moving objects, relying primarily on the scanner itself to measure self-motion. An extended Kalman filter with a single constant velocity model is used in a multi-hypothesis tracker. As opposed to our work, their emphasis appears to be on mapping in the presence of moving objects, rather than the real-time detection of moving objects when no map is needed.

Using a map, such as an occupancy grid, appears to offer a convenient way of detecting moving objects by simply observing the changes in occupancy values for each location. However, maintaining an occupancy grid is expensive; (Lindstrom and Eklundh, 2001) addressed this problem with a sparse representation of open space. Yet, the grid does not solve the shape-change problem because we cannot disregard the possibility that an object was there already but could not be detected due to occlusion or because it was out of range. The effect of range limits is particularly intractable because it depends on the unknown target reflectivity.

Several papers describe indoor people tracking systems that use laser scanners. Shape-change effects are mild when tracking people because people are compact compared to typical sensing ranges and do not have flat surfaces. Although a moving scanner will see shape change in large objects such as desks, large objects can simply be discarded because they are clearly not people.

In (Fod et al., 2002), motion is measured by registering old and new scans using chamfer fitting. A constant velocity, constant angular velocity Kalman filter is used. Because the scanner is placed above the leg level, a rigid body model is satisfactory. Although this paper does not use moving scanners, it is noteworthy because of its attempt to quantitatively evaluate performance without ground truth by measuring the position noise of stationary tracks, the measurement residue of moving tracks, and the occurrence of false positive and false negative errors in moving object detection.

There is a large body of literature on tracking techniques developed for long-range radar which can be applied to robotic applications (Bar-Shalom and Fortmann, 1988). However, the low resolution and long

range means that all objects are treated as points. In our problem, we deal with objects which are closer to the sensor. Consequently, a single object is treated as a collection of points, rather than a single one.

3. SYSTEM DESCRIPTION

In this section, the main steps of the algorithm used in our system are described. A more detailed description of these steps is presented in (MacLachlan and Mertz, 2006).

3.1 Detection, Tracking, and Prediction

Objects in the vicinity of the vehicle are detected using measurements collected by the vehicle's LADAR line scanners. The objects might be static objects in the environment, or moving objects (see Figure 2 for an overall depiction). The detection and tracking process involves the execution of the following sequence of tasks: object detection, object tracking, and prediction.

Object Detection. The first step of the algorithm is the grouping of the 3-D points measured by the LADAR into potential objects. The objects might be static objects in the environments (which are discarded later), or moving objects. The algorithm can handle people or vehicles as moving objects. Each LADAR scan is segmented into objects. Each object is summarized by a corner or a line which is fitted to the set of points belonging to the object. The corner- and end-points are the feature points used for describing the object.

Object Tracking. The objects detected in the current scan are then matched with segments from previous scans. A measure of match quality is extracted for each pair of objects. If they do match, they are considered the same objects and the motion of the feature points calculated from the match is then fed into a Kalman filter which calculates the velocity of the object. The output of the detection and tracking algorithm is a set of objects and their attributes (i.e., position and velocity). To handle the shape change problem, we apply a separate track validation procedure that determine whether recent observations are consistent with rigid body motion under the dynamic model, and whether there is sufficient evidence to conclude that the object is definitely moving. To address the problem of tracking objects over a 360° envelope around the vehicle, the detection and tracking algorithm are executed over four sensors arranged around the vehicle with overlapping fields of view. As a result, it is necessary to “hand off” objects tracked in one field of view to the next. The fusion of the four sensors happens at the object level. The segmentation and feature extraction is done for each sensor scan separately. There is only one object list and each scan updates the objects within its own field-of-view. Objects which are seen by two sensors are updated twice per cycle.




Fig. 2. Overview of the detection and tracking system . At the center of the figure is the vehicle (viewed from above), which carries four LADARs (one on each side).

Prediction. The last part of the system is the generation of the predicted trajectories from the objects detected and tracked in the fields of view. We assume that the trajectories that a given object may choose in the future follow a given probability distribution which may not be normal or even unimodal. In that case, it becomes impossible to represent the distribution of trajectories parametrically (e.g., by its mean and variance) and it becomes important to resort to non-parametric techniques. In the current approach, representative trajectories are sampled according to the underlying probability distribution. Samples tend to concentrate in areas of the space of object trajectories that are more likely, while they tend to be sparse in areas that are unlikely. This component use a particle filter to predict object positions seconds into the future using only the current motion estimate. This approach enables the use of more sophisticated prediction models. For example, it can support multi-modal distributions of trajectories that cannot be represented by, for example, a simple dynamic model from a Kalman filter.

The predictions have been used for collision warning (MacLachlan and Mertz, 2006). More precisely, the system generates the probability of collision for each proposed vehicle's trajectory at varying time horizons. These predicted trajectories can also be used as input to a vehicle's planner to implement avoidance of dynamic obstacles.

The algorithms described above are fast enough to keep up with a scanner acquisition rate of 75 Hz when running on a 600 MHz embedded processor.

3.2 Parameters

There are multiple parameters that affect the operation of the tracker. These have been empirically tuned for our particular scanner and application. We summarize the most relevant here:

- 1) A track is considered apparently moving if it has been tracked for at least 15 cycles, and the speed is greater than 0.75 m/s.
- 2) A track is valid when it is apparently moving, the standard deviation of its velocity estimate is less than 0.8 m/s, and has maintained a history of consistency for at least 10 cycles.
- 3) A minimum of 3 points are required to create a new object track; a minimum of 2 points are required to keep the track alive.
- 4) To keep the computational load low, range measurements longer than 40 m are ignored.

4. EXPERIMENTAL SETUP

4.1 Metrics

The performance of the system can be evaluated along many different axes, each with different metrics. Accordingly, we defined a number of metrics and corresponding experiments and carried them out using one of the experimental setups described before. The rest of this section describes the key experimental results according to these metrics, which are:

Detection Distance: This is the distance to the object at the time the track is considered valid.

Velocity Error: This is the difference between the velocity measured by the system and the ground-truth velocity. Mean and standard deviation of velocity error are reported.

Velocity Delay: This is the delay between the time at which an object is detected and the time at which its velocity is estimated. As opposed to the position measurement process, where no target dynamic model is used and from which estimates are immediately available, a Kalman filter is used to compute target velocity estimates. The initial velocity of an object is assumed to be zero; it takes several cycles to establish a reliable velocity estimate. However, there is a tradeoff between the accuracy of the velocity estimate and the delay in acquiring it. An attempt to obtain a valid estimate faster implies a relaxation of the uncertainty allowed for that estimate to be considered valid. In our system, there has to be a consistency in the history of the track before an estimate can be declared well grounded. This consistency is evaluated using multiple criteria. For example, an estimate has to undergo a minimum number of cycles; the standard deviation of the estimate should not exceed a maximum threshold, and the estimate should remain consistent for at least a certain minimum time. For our experiments, a track must have data associated for at least 15 iterations (equivalent to 0.4 seconds at 37.5 Hz). Similarly, the maximum standard deviation of the velocity estimate allowed in a track to be reported as valid is 0.8 m/s. Finally, the velocity estimate should remain




Fig. 3. Estimation of target velocity. A pedestrian walking at a constant speed of 2 m/s is tracked. The system reports a valid velocity estimate after 0.8 seconds, as shown in the top figure. The standard deviation of the velocity estimate, plotted in the bottom figure, is one of several criteria used to validate the track.

consistent (i.e., without significant variations of its standard deviation) for at least 10 iterations.

The application of these criteria is illustrated in Figure 3, which shows the velocity measurement of a pedestrian walking at a constant speed in front of a static vehicle (NavLab11), and the corresponding uncertainty reported by the Kalman filter. As the person is detected, the system starts estimating its velocity (top plot). After 0.4 s (equivalent to 15 iterations), the first criterion is satisfied. As the velocity estimate converges to the true value, its standard deviation decreases below the 0.8 m/s threshold at 0.46 s, as shown in the bottom plot. This satisfies the second requirement. The standard deviation continues to decrease, and eventually reaches a steady-state value. As shown in the top plot, a consistent velocity estimate (third requirement) is produced after 10 iterations producing, a valid velocity estimate is produced at 0.8 s, and the track is considered valid.

Track Breakup: The position estimation process can be negatively affected by several causes. The system fails to detect a target when the target is occluded, when it has poor reflectivity at the infrared frequency at which the scanner operates, or when objects are too close to each other and it is not clear whether to segment the data as one or more objects. The latter is known as clutter, and can cause the spontaneous disappearance of tracks when a target moves close to another object, even though it is not visually occluded from the scanner.

The system continuously collects measurements and seeks to establish relationships among groups of adjoining points to determine whether they belong to the same object. This determination may fail for a number of reasons, including occlusion from background clutter, poor reflectivity, or objects in close proximity to each other (in which case it is not clear whether to segment the

data as one or more objects. When this occurs, a single target is tagged with many different labels as it is being tracked. This has a negative impact in the velocity estimation, since every time a new target is detected, there is a time delay until a new target velocity fix is available, as described before. The re-labeling of the target, also known as *track breakup*, results in an accumulation of time during which the target is not accurately tracked.

Prediction: One way to assess the performance of the approach is to measure the rate of correctly anticipated potential collisions with the vehicle. This is a difficult metric to evaluate because of its subjective nature (the only way to get real “ground truth” is to actually collide with the object, a procedure that is not practical when the moving objects of interest are humans!)

4.2 Experimental Platforms

We used several testing platforms during the design of this system. We have tested our algorithms using data collected from these four configurations:

1. Tabletop: The laser scanners are on a fixed platform. This configuration is useful to characterize the baseline performance of the sensors and of the system on a completely stationary platform (i.e., without even engine vibrations or other effects from a “live” vehicle).

2. Demo III XUV (Shoemaker and Bornstein, 1998): Four scanners are mounted on the XUV. Data was collected from natural environments in central Pennsylvania and northwestern Maryland. This configuration is used to validate the performance and operation of the system on the target platform.

3. NavLab11: The CMU test vehicle is a Jeep Wrangler with three scanners, one in front and one on each side. It was driven at various speeds on and off-road, taking data in normal traffic and under controlled circumstances. This platform is particularly valuable for evaluating the performance of the system at high speed (e.g., 20 mph or higher).

4. Transit vehicles: As part of different, but related, project, we mounted two scanners on two transit vehicles, one scanner on each side (MacLachlan and Mertz, 2006). The predictive obstacle detection and tracking was part of a side collision warning system. We collected hundreds of hours of data during normal operation. The data was used to calibrate and evaluate the system. We use some of this data in this report since it is the largest data set ever collected on the use of detection and tracking systems in an uncontrolled environment. This data provides valuable information in addition to the controlled experiments conducted on XUV or Navlab11. Also, this system provided invaluable lessons that guided the design of the system described in this report.

The laser scanners have an update rate of 75 Hz or 37.5 Hz, depending on the resolution of 1° or 0.5°. Importantly, this processing rate implies that the

magnitude of the objects’ motion at each cycle is very small, thus facilitating the tracking.

It is important to note that we use these commercial off-the-shelf sensors for convenience of experimentation, but other sensors can also be used with this approach (we have successfully tested our system using 3-D points collected from a mobility LADAR). All the quantitative results presented in the report are relative to the performance of these sensors, but the algorithms are for the most part independent of the sensors. In particular, more recent versions of the scanners allow for finer angular resolution, which is a major limitation of the current implementation.

5. EXPERIMENTAL EVALUATION

5.1 Base line performance

Evaluation of tracking systems is difficult, since it is hard to provide target ground truth. A formal assessment of such systems in vehicular applications is rarely found in the literature. For this reason, we conducted a number of experiments using a small mobile robot with all-terrain driving capabilities (Fig. 1). We took advantage of its controlled motion capabilities to establish the baseline performance of the system. The robot is equipped with wheel encoders, a fiber-optic gyro (yaw), and a laser rangefinder. These sensors provide accurate position and pose estimates, which were used to establish the ground truth for evaluating the tracking system. Besides, this robot was primarily used as a target during high-speed tests for “pedestrian” detection with the remotely controlled XUV. (Due to safety concerns one can not perform these experiments with humans.)

We conducted experiments using the small robot with both the XUV and NavLab11¹, for a combined total of 12 experimental runs. The robot was set in motion at a constant speed, and the vehicle collected data while maneuvering around the robot. The robot motion information was then compared with the estimates reported by the tracking system. Some of these results, obtained at vehicular speeds of 16 and 18 mph, are summarized in Table 1.

Vehicle velocity	16 mph	18 mph
Mean error	0.089 m/s	-0.0728 m/s
Std. dev. error	3.98 cm/s	6.93 cm/s
Velocity delay	1.1 s	2.0 s
Track duration	11.5 s	6.5 s
Target velocity	0.902 m/s	0.883 m/s
Detection distance	25.5 m	36.48 m
Target direction	Same as vehicle	Toward vehicle

Table 1. Ground truth experiments using a controlled target

¹ In all of the experiments conducted using NavLab11, including tests at higher speeds, for safety reasons the vehicle was manually driven with no computer interfering with the driving.

5.2 Velocity Delay (with High Vehicle Speed):

Since the robot does not imitate the human gait, we also performed controlled experiments involving pedestrians moving at a constant velocity, while the vehicle (NavLab11) is moving at a higher speed. These experiments were performed in an open environment (a quiet street nearby a park), with moderate presence of clutter. Both the pedestrians (targets) and the vehicle were moving on flat ground. All pedestrians carried a stop watch each, and used it to determine their traveling speed, based on distance markings drawn on the sidewalk. We collected data from a total of 48 individuals. The results are summarized in Table 2. As shown, the system is able to consistently report accurate velocity estimates while NavLab11 travels at speeds of up to 40 mph. Pedestrians were initially detected from as far as 36.88 m, which is close to the maximum range of 40 m.

The amount of clutter in these experiments can be appreciated in Figure 4, which illustrates one experimental run of NavLab11 driving at 25 mph. In this run, the pedestrian was never significantly affected by clutter. No breakups occurred during this experiment.

The delay in velocity estimation is significant in these experiments. It should be noted that this delay affects only the reporting of an accurate target velocity by the system and that a much shorter delay can be used to report the detection of a moving object if one accepts to sacrifice accuracy of the velocity estimate. The parameters used here are conservative and are designed to guarantee a standard deviation of estimated velocity lower than 0.1m/s.

5.3 Track Breakup

To assess the effect of track breakup in typical environments, we conducted a series of 5 experiments in which people were walking alongside an XUV (manually guided using a pendant) in an off-road environment. We describe in detail one representative experiment in which four people (referred to as “targets” from now on) were moving around the XUV, traversing a distance of 92.2 m at relatively constant speed. The vehicle speed varied between 0.9 and 1.3 m/s. The test was conducted outdoors, in a rural environment. Table 3 summarizes track breakup occurrence. In this experimental run, target A moved always ahead of the vehicle, while periodically crossing from one side to another. Targets B and C always remained behind and close to the XUV (less than




Fig. 4. A pedestrian, identified as Track 2340, moves at constant velocity. NavLab11 estimates the pedestrian's velocity while driving at 25 mph. The red line indicates the pedestrian's estimated velocity and direction. Other objects in the scene are identified by rectangular boxes. Raw scanner measurements appear as blue dots.

4 m), and were never occluded nor significantly affected by clutter. Target D followed the XUV from slightly farther away and eventually walked across tall grass areas, to the point of being lost in the clutter for extended amounts of time. As shown in the table, the system performed well, reporting valid velocity estimates 95.99%, 79.49%, and 96.37% of the time for targets A, B, and C, respectively. Similarly, there were few breakups for these three targets, being as low as 4 for target A, and as high as 10 for target B.

Target D was frequently occluded or cluttered by the tall grass and suffered as many as 57 breakups, which precluded the computation of valid estimates more than 51% of the time. At some point, the system assigned a new track for this target every 0.1 s, since the target walked too close to a patch of tall grass, even though the scanners had an unobstructed view of it.

5.4 Prediction

We include here data acquired with a version of the system that was used on transit vehicles. In this case, a warning is issued whenever the predicted trajectory of an object intersects the predicted trajectory of the vehicle. More precisely, a warning is issued whenever the probability of a collision rises above a certain threshold. Two level of warnings, an “alert” and an “imminent warning” for different degrees of danger are generated.

<i>Vehicle speed, m/s (mph)</i>	<i>Target velocity, m/s</i>	<i>Estimated velocity, m/s</i>	<i>Mean velocity estimation error, m/s</i>	<i>Target detection distance, m</i>	<i>Velocity delay, s</i>	<i>Standard deviation of velocity, m/s</i>
9.1 (20.36)	1.62	1.686	-0.0661	34.7	1.1	0.046
10.5 (23.49)	2.18	2.1	0.077	36.88	1.8	0.0483
11 (24.6)	3.95	4.08	-0.132	34.86	3.5	0.0656
14.2 (32)	1.71	1.78	0.061	33.73	1.4	0.075
17.7 (40)	2.89	2.92	0.027	34.26	1.42	0.081

Table 2. Tracking pedestrians at high speed.

Target	No. of track breakups	Average velocity delay (s)	Minimum velocity delay (s)	Maximum velocity delay (s)	Percentage of time with valid velocity estimate
A	4	2.1	1.5	2.8	95.99 %
B	10	3.57	0.2	10.7	79.49 %
C	9	0.81	0.2	2.9	96.37 %
D	57	4.95	0.8	15.0	28 %

Table 3. Track breakup analysis: four pedestrians walking alongside an XUV moving at low speed.

The warnings are based on computing, at each cycle, the most likely trajectory for each detected moving target and to intersect it with the predicted trajectory of the vehicle (in this particular experiment, the vehicle was manually driven by an independent driver). Predicting trajectories of detected objects by taking into account typical behavior is in fact a challenging problem and so is the evaluation of the prediction algorithm.

For the purpose of documenting performance of the detection and tracking system in the context of an overall safety system, we analyzed the warnings issued over 5 hours worth of data collected in urban environments. Although this analysis is far more qualitative than the other results presented in this report, it is important because it uses one of the few datasets that was acquired in an unbiased, uncontrolled manner, i.e., we had no control over the environment, the motion of the people and vehicles, and the motion of the vehicle, which was driven by an independent driver. For each warning that was issued we determined if it was a true, i.e. correct, warning. We determined the reason of all the false warnings. Table 4 shows the absolute number of warnings, the relative number for each category (percentage of the total number of warnings) for each cause, and the warning rate, for the left and right sides.

The most common situations that cause true warnings are vehicles passing and fixed objects in the path of a turning vehicle. On the right side there are additional true warnings caused by pedestrians entering the vehicle or walking towards it when the vehicle has not yet come to a full stop. These are counted as “false positive” in this particular scenario, but they would be true positive in a scenario in which people approach the vehicle from any

direction are considered threats.

A majority of the alerts are true alerts, whereas a majority of the imminent warnings are false positives. The most common reason for false imminent warnings is that the velocity was incorrect, but as explained below this kind of error is not very serious. The main sources of errors are:

Vegetation: The warning is triggered by vegetation (grass, bush, etc.). The system performs correctly, but the warning is regarded as a nuisance because grass or bushes are not considered dangerous. For an autonomous vehicle, these warnings can be eliminated by integrating the safety system with other terrain classification components.

False velocity: These are the outliers in the velocity measurement discussed earlier. The velocity estimate is sometimes slightly off, which increases the probability enough to cross the warning threshold. This error is not extremely serious, because it is only an error in the degree of danger.

No velocity: These are the delay in velocity measurements discussed earlier. An object is detected but the velocity is not yet established and is therefore assumed to be zero. This leads to false warnings when the vehicle approaches another vehicle with similar speed.

The error rates listed in Table 4 are only the cases where a warning was issued when there should not have been one (false positive warnings). Many of the reasons mentioned above could also cause false negative warnings, i.e. missed warnings. The rate of false negative warnings is very hard to determine, because one has to look through all the data to find situations where a warning should have been given. What we did instead is to stage collisions and determined if a warning was

	absolute				relative [%]				rate [1/hour]			
	alert		imminent		alert		imminent		alert		imminent	
	right	left	right	left	right	left	right	left	right	left	right	left
True	60	94	15	9	59	71	47	26	12.0	18.8	3.0	1.8
Vegetation	10	3	2	0	10	2	6	0	2.0	0.6	0.4	0.0
false velocity	21	28	10	20	21	21	31	57	4.2	5.6	2.0	4.0
no velocity	0	2	1	0	0	2	3	0	0.0	0.4	0.2	0.0
ground return	10	4	3	3	10	3	9	9	2.0	0.8	0.6	0.6
Other	1	2	1	3	1	2	3	9	0.2	0.4	0.2	0.6
Sum	102	133	32	35	100	100	100	100	20.4	26.6	6.4	7.0

Table 4. True and false positive warnings.

missed for those situations.

We staged 30 collisions or near collisions. 17 of those were used to calibrate the system and the remaining 13 were the test set. In all thirteen cases the system gave correct warnings. Since no false warnings were observed, we can only give an upper bound on the false warnings rate: With a 90% confidence level the false warning rate for these scenarios is less than 0.16.

CONCLUSIONS

We have described the design and experimental evaluation of a predictive mover detection and tracking system, capable of operating from a moving vehicle in real-time. In our approach, the detection-tracking-prediction elements are integrated into a single system.

The system's base line performance was evaluated by conducting experiments using a small mobile robot as a controlled target to provide ground truth. Since the robot does not imitate the human gait, we also performed tests with humans using NavLab11 and a Demo III XUV.

The system has proven capable of detecting humans moving as fast as 4 m/s at distances up to 38 m, from a vehicle moving at speeds as high as 40 mph, and measured the target's velocity with an error as small as 0.061 m/s. The prediction capabilities were tested using data collected in urban environments (performance is summarized in the previous section).

The experiments have shown that the integrated approach described in this document can be used in a system that can detect and track objects and predict the trajectories of objects and the corresponding probabilities of collision with the vehicle. The approach has still many limitations. Areas which can still be improved are: decrease in velocity delay and number of track breakups, especially near clutter, and filtering of vegetation. In addition, current results show that the predictions computed from the output of the detection and tracking system can be used effectively to predict possible collision with future vehicle paths, thus motivating further development of the prediction system.

ACKNOWLEDGEMENTS

This work was conducted through collaborative participation in the Robotics Consortium sponsored by the U. S. Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-209912. We would like to thank General Dynamics Robotic Systems for their support.

REFERENCES

- Bar-Shalom, Y., and Fortmann, T. E., 1988: Tracking and Data Association, Mathematics in Science and Engineering, Academic Press.
- Fod, A., Howard, A., and Mataric, M., 2002: A Laser-Based People Tracker, Proc. IEEE International Conference on Robotics and Automation, pp. 3024-3029.
- Fuerstenberg, K., Dietmayer, K., and Willhoeft, V., 2002: Pedestrian Recognition in Urban Traffic Using a Vehicle Based Multilayer Laserscanner, Intelligent Vehicle Symposium, pp. 31-35.
- Fuerstenberg, K., Scholz, J., 2005: Reliable Pedestrian Protection Using Laserscanners, Proceedings 2005 IEEE Intelligent Vehicles Symposium, 6-8 June 2005, pp. 142 – 146.
- Hebert, M., 2000: Active and Passive Range Sensing for Robotics, Proc. IEEE Int. Conf. Robotics and Automation, vol. 1, pp. 102-110.
- Lindstrom, M., and Eklundh, J., 2001: Detecting and Tracking Moving Objects From a Mobile Platform Using a Laser Range Scanner, IEEE International Conference on Intelligent Robots and Systems, pp. 1364-1369.
- MacLachlan, R., and Mertz, C., 2006: Tracking of Moving Objects from a Moving Vehicle Using a Scanning Laser Rangefinder, Proc. 9th International IEEE Conference on Intelligent Transportation Systems, September 17-20, Toronto, Canada.
- Shoemaker, C.M., and Bornstein, J. A., 1998: The Demo III UGV Program: a Testbed for Autonomous Navigation Research, Proceedings of the IEEE International Symposium on Intelligent Control, Gaithersburg, MD, September 1998, pp. 644-651.
- Sun, Z., Bebis, G., and Miller, R. 2006: On-road Vehicle Detection: a Review, IEEE Transactions on Pattern Analysis and Machine Intelligence, Volume 28, Issue 5, May 2006. pp. 694 – 711.
- Streller, D., Fuerstenberg, K., and Dietmayer, K., 2002: Vehicle and Object Models for Robust Tracking in Traffic Scenes Using Laser Range Images, Intelligent Transportation Systems, 2002, pp. 118-123.
- Wang, C., Thorpe, C., and Thrun, S., 2003: Online Simultaneous Localization and Mapping with Detection and Tracking of Moving Objects: Theory and Results From a Ground Vehicle in Crowded Urban Areas, Proc. IEEE International Conference on Robotics and Automation, pp. 842-849.